

AD-A055 045

MASSACHUSETTS INST OF TECH CAMBRIDGE ARTIFICIAL INTE--ETC F/G 6/4
REPRESENTING VISUAL INFORMATION, (U)

MAY 77 D MARR
AI-M-415

UNCLASSIFIED


N00014-75-C-0643

NL

1 OF 1
AD
A055045



END
DATE
FILED
7-78
DDC



AD No.

DDC FILE COPY

ADA055045

FOR FURTHER TRAN

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE			READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 14	2. GOVT ACCESSION NO. ADM-415	3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle) 6 REPRESENTING VISUAL INFORMATION		5. TYPE OF REPORT & PERIOD COVERED MEMO	
7. AUTHOR(s) D. Marr		6. PERFORMING ORG. REPORT NUMBER 15 N00014-75-C-0643	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Artificial Intelligence Laboratory 545 Technology Square Cambridge, Massachusetts 02139		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Blvd Arlington, Virginia 22209		12. REPORT DATE 11 May 77	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research Information Systems Arlington, Virginia 22217		13. NUMBER OF PAGES 1264p. 62	
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this document is unlimited.			
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) <i>This document has been approved for public release and sale; its distribution is unlimited.</i>			
18. SUPPLEMENTARY NOTES None			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) vision, image processing, surface representation, 3-D vision, review article, motion, stereopsis			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) Vision is the construction of efficient symbolic descriptions from images of the world. An important aspect of vision is the choice of representations for the different kind of information in a visual scene. In the early stages of the analysis of an image, the representations used depend more on what it is possible to compute from an image than on what is ultimately desirable, but later representations can be more sensitive to the specific			

BLOCK 20 CONCLUDED

needs of recognition. This essay surveys recent work in vision at M.I.T. from a perspective in which the representational problems assume a primary importance. An overall framework is suggested for visual information processing, in which the analysis proceeds through three representations; (1) the primal sketch, which makes explicit the intensity changes and local two-dimensional geometry of an image, (2) the 2½-D sketch, which is a viewer-centered representation of the depth, orientation and discontinuities of the visible surfaces, and (3) the 3-D model representation, which allows an object-centered description of the three-dimensional structure and organization of a viewed shape. Recent results concerning processes for constructing and maintaining these representations are summarized and discussed.

**MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY**

AIM 415

May 1977

REPRESENTING VISUAL INFORMATION

D. Marr

SUMMARY: Vision is the construction of efficient symbolic descriptions from images of the world. An important aspect of vision is the choice of representations for the different kinds of information in a visual scene. In the early stages of the analysis of an image, the representations used depend more on what it is possible to compute from an image than on what is ultimately desirable, but later representations can be more sensitive to the specific needs of recognition. This essay surveys recent work in vision at M.I.T. from a perspective in which the representational problems assume a primary importance. An overall framework is suggested for visual information processing, in which the analysis proceeds through three representations: (1) the primal sketch, which makes explicit the intensity changes and local two-dimensional geometry of an image, (2) the $2\frac{1}{2}$ -D sketch, which is a viewer-centered representation of the depth, orientation and discontinuities of the visible surfaces, and (3) the 3-D model representation, which allows an object-centered description of the three-dimensional structure and organization of a viewed shape. Recent results concerning processes for constructing and maintaining these representations are summarized and discussed.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643.

Contents

0 INTRODUCTION

- 0.1 Understanding information processing problems
- 0.2 Understanding vision

1 EARLY PROCESSING PROBLEMS

- 1.0 The primal sketch
- 1.1 Random-dot interference patterns
- 1.2 Texture discrimination
- 1.3 Discrimination ability
- 1.4 Light source effects
- 1.5 Regions from a discriminant
- 1.6 Lightness

2 PROCESS-ORIENTED THEORIES

- 2.0 Introduction
- 2.1 Structure from motion
- 2.2 Stereopsis

3 INTERMEDIATE PROCESSING PROBLEMS

- 3.0 Introduction
- 3.1 Difficulties with idea of image segmentation
- 3.2 Reformulating the problem
- 3.3 General classification of shape representations
- 3.4 Some observations from psychophysics
- 3.5 The computational problem
- 3.6 A possible form for the $2\frac{1}{2}$ -D sketch

4 LATER PROCESSING PROBLEMS

- 4.0 Introduction
- 4.1 Nature of an object-centered coordinate system
- 4.2 Shapes having natural coordinate systems
- 4.3 Finding the natural coordinate system from an image
- 4.4 Interpreting the image of a single generalized cone
- 4.5 Surfaces composed of two or more generalized cones
- 4.6 Some comments on the limitations of this theory

5 DISCUSSION

6 REFERENCES



O: Introduction

0.1: Understanding information processing tasks

Vision is an information processing task, and like any other, it needs understanding at two levels. The first, which I call the computational theory of an information processing task, is concerned with what is being computed and why; and the second level, that at which particular algorithms are designed, with how the computation is to be carried out (Marr & Poggio 1977a). For example, the theory of the Fourier transform is a level 1 theory, and is expressed independently of ways of obtaining it (algorithms like the Fast Fourier Transform, or the parallel algorithms of coherent optics) that lie at level 2. Chomsky calls level 1 theories competence theories, and level 2 theories performance theories. The theory of a computation must precede the design of algorithms for carrying it out, because one cannot seriously contemplate designing an algorithm or a program until one knows precisely what it is meant to be doing.

I believe this point is worth emphasizing, because it is important to be clear about the level at which one is pursuing one's studies. For example, there has recently been much interest in so-called cooperative algorithms (Marr & Poggio 1976) or relaxation labelling (Rosenfeld, Hummel & Zucker 1978). The attraction of this technique is that it allows one to write plausible constraints directly into an algorithm, but one must remember that such techniques amount to no more than a style of programming, and they lie at the second of the two levels. They have nothing to do with the theory of vision, whose business it is to derive the constraints and characterize the solutions that are consistent with them.

0.2: Understanding vision

If one accepts in broad terms this statement of what it means to understand an information processing task, one can go on to ask about the particular theories that one needs to understand vision. Vision can be thought of as a process, that produces from images of the external world a description that is useful to the viewer and not cluttered by irrelevant information. These descriptions, in turn, are built or assembled from many different but fixed representations, each capturing some aspect of the visual scene. In this article, I shall try to present a summary of our work on vision at M.I.T. seen from a perspective in which the representational problems assume a primary importance. I shall include summaries of our present ideas as well as of completed work.

The important point about a representation is that it makes certain information explicit (cf the principle of explicit naming, Marr 1976). For example, at some point in the analysis of an image, the intensity changes present there need to be made explicit, so does the geometry — of the image and of the viewed shape — and so do other parameters like color, motion, position and binocular disparity. To understand vision thus requires that we first have some idea of which

representations to use, and then we can proceed to analyze the computational problems that arise in obtaining and manipulating each representation. Clearly the choice of representation is crucial in any given instance, for an inappropriate choice can lead to unwieldy and inefficient computations. Fortunately, the human visual system offers a good example of an efficient vision processor, and therefore provides important clues to the representations that are most appropriate and likely to yield successful solutions.



This point of view places the nature of the representations at the center of attention, but it is important to remember that the limitations on the processes that create and use these representations are an important factor in determining their structure, because one of the constraints on vision is that the description ultimately produced be derivable from images. In general, the structure of a representation is determined at the lower levels mostly by what it is possible to compute, whereas later on they can afford to be influenced by what it is desirable to compute for the purposes of recognition.

1: Early processing problems


1.0: The primal sketch

There are two important kinds of information contained in an intensity array, the intensity changes present there, and the local geometry of the image. The primal sketch (Marr 1976) is a primitive representation that allows this information to be made explicit. Following the clues available from neurophysiology (Hubel & Wiesel 1962), intensity changes are represented by blobs and by oriented elements that specify a position, a contrast, a spatial extent associated with the intensity change, a weak characterization of the type of intensity change involved, and a specification of points at which intensity changes cease (so-called termination points). The representation of local geometry makes explicit two-dimensional geometrical relations between significant items in an image. These include parallel relationships between nearby edges, and the relative positions and orientations of significant places in the image. These significant places are marked by "place-tokens", and they are defined in a variety of ways, by blobs or local patches of different intensity, by small lines, and by the ends of lines or bars. The local geometrical relations between place-tokens are represented by inserting virtual lines that join nearby place-tokens, thus making explicit the existence of a relation between the two tokens, their relative orientation, and the distance between them (Marr 1976 figure 12a).


I. The primal sketch makes explicit information held in an intensity array (Ia). There are two kinds, one concerns the changes in intensity, and this is represented by oriented edge, line and bar elements, associated with which is a measure of the contrast and spatial extent of the intensity change. The other kind of information is the local two-dimensional geometry of significant places in the image. Such places are marked by "place-tokens", which can be defined in a variety of ways, and the geometric relations between them are represented by inserting "virtual lines" between nearby tokens. (Marr 1976 figures 7 and 12a).




a.



b.



c.



d.



2. 2a and 2c are random-dot interference patterns of the kind described by Glass (1969). 2b and 2d exhibit the results of running the algorithm described in the text and figure 3. The neighborhood radius was such that roughly 8 neighbors were included. (Stevens 1977 figure 5).

1.1: Random-dot interference patterns

The idea of place-tokens and of this way of representing geometrical relations arose from considering the computational problems that are posed by early visual processing, and one of the questions we have been asking is, can one find any psychophysical evidence that the human visual system makes use of a similar representation? We have recently obtained two results related to this point. Stevens (1977) has examined the perception of random-dot interference patterns (figure 2), constructed by superimposing two copies of a random dot pattern where one copy has undergone some composition of expansion, translation, or rotation transformations (Glass 1969). He found that a simple algorithm suffices to account quantitatively for human performance on these patterns. The algorithm consists of three steps:

(1) Each dot defines a place-token. For example some dots can be replaced by small lines or larger blobs without disrupting the subjective impression of flow.

(2) Virtual lines are inserted between nearby place-tokens, and the neighborhood in which the virtual lines are inserted depends in a predictable way on the density of the dots.



(3) The orientations of the virtual lines attached to all the points in each neighborhood are histogrammed, and locally parallel organization is found by searching for a peak in this histogram. The bucket width that best matches human performance is about 10 degrees.


The details of these steps are set out in figure 3. The interesting features of the algorithm are; (a) It is not iterative. Stevens could find no evidence that human performance rests on a cooperative algorithm, although this type of problem is ideal for that approach. (b) The algorithm is purely local. No global-to-local or top-down interactions are necessary to explain human performance. (c) What the algorithm finds is locally parallel organization. In this case, the organization lies in the virtual lines constructed between nearby dots, but locally parallel organization among the real edges and lines in an image also forms an important part of the structure of an image (Marr 1978).

1.2: Texture discrimination

The second study is one by Schatz (1977) on texture vision discrimination. Marr (1978) suggested that such discriminations could be carried out by first-order discriminations acting on the description in the primal sketch (p.50). Marr supposed that certain grouping processes were needed before the discriminations are made in order to account for the full range of human texture discrimination, but in a careful examination of the problem, Schatz found that many of the examples he constructed could be explained by assuming that the discriminations are made only on real edges or on virtual lines inserted between neighboring place-tokens. If this were generally true, it would stand in elegant relation to Julesz's (1975) conjecture, that a necessary condition for the discriminability of two textures is that their dipole statistics differ. This condition is known not to be sufficient, a state of affairs that one can view as implying that we

3. The algorithm for computing locally parallel structure has three fundamental steps. Place tokens that are defined in the image are the input to the algorithm, which is applied in parallel to each one. Since, in the case of the Moire dot patterns, each dot contributes a place token, the first step is to construct a virtual line from that dot to each neighboring dot (within some neighborhood centered on the dot). A virtual line represents the position, separation, and orientation between a pair of neighboring dots. To favor relatively nearer neighbors, relatively short virtual lines are emphasized. The second step is to histogram the orientations of the virtual lines that were constructed for each of the neighbors. For example, the neighbor D would contribute orientations AD, DF, DG, and DH to the histogram. The final step (after smoothing the histogram) is to determine the orientation at which the histogram peaks, and to select that virtual line (AB) closest to that orientation as the solution. (Stevens 1977 figure 4).





4. The pattern 4a contains two regions, one of whose line segments has the orientation distribution shown in 4b, and the other has the distribution 4c. Surprisingly, three orientations cannot be distinguished from a random orientation distribution. If human texture discrimination is based on first-order discriminations acting on the description held in the primal sketch, the discriminants that can be brought to bear on this information are weak.

have access to only a proper subset of all dipole statistics. It is possible that this proper subset consists only of real edges and of the virtual lines that join nearby place-tokens.

1.3: Discrimination ability

If one accepts that texture discrimination relies upon first-order discriminations of this type, it is natural to ask how sensitive are the particular discrimination functions that we can bring to bear on an image. Riley (1977) has found evidence that the available functions are extremely coarse. For example, figure 4 consists of a background in which the fine segments have a random orientation, surrounding a square containing lines of only three orientations. Surprisingly, the square cannot be discerned without scrutiny. One interpretation of this and related findings is, that discriminations on orientations other than horizontal and vertical are made on the output of 5 channels, each nearly binary, and with an angular width of about 35 degrees -- in other words, only very little information is available about the distribution of orientations in an image. It appears that our discrimination ability is as poor or poorer for the other stimulus dimensions, for example intensity distribution (Riley 1977).

1.4: Light source effects

In another study concerned with what can be extracted from an image, Ullman (1976a) enquired about the possible physical basis for the subjective quality of fluorescence, which is normally associated with the presence of a light source. He noted that at a light source boundary, the ratio of intensity to intensity gradient changes sharply, whereas this is not true at reflectance boundaries unless the surface orientation changes sharply. He showed that, in the mini-world of Mondrians, the discriminant to which this leads predicts human performance satisfactorily.

K. Forbus (in preparation) has extended this work to the detection of surface luster. Since glossiness is due to the specular component of a surface reflectivity function, one can treat the detection of gloss as essentially the detection of light sources that appear reflected in a surface (see Beck 1974), and this depends ultimately on the ability to detect light sources. Forbus divided the problem into three categories; (a) in which the specularity is too small to allow gradient measurements, (b) in which both intensity and gradient measurements are available, but the specularity is local (as it is for a curved surface or a point source), and (c) in which the surface is planar and the source is extended. He derived diagnostic criteria for each case.

1.5: Regions from a discriminant

Whenever a region is defined in an image by a predicate, for example by a difference in texture or brightness, one faces the problem of delimiting the region accurately. There are two approaches to designing algorithms for this problem; one is to use the predicate

directly, deciding whether a given location lies within or without the region by testing some function of the predicate there. The second approach is to differentiate the predicate, defining the region by its boundaries rather than by properties of its interior.





The difficulties with the problem arise because one is usually ignorant beforehand of the scale at which significant predicate signals may be gathered. For example, suppose one wished to find the boundary between two regions that are distinguished by different densities of dots. Dot density has to be measured by selecting a neighborhood size and counting the number of dots that lie within it. If the neighborhood size is too large, one may not be able to resolve the regions. If it is so small as to contain zero, one or two dots, natural fluctuations may obscure any changes in density.

One solution to this problem is to make the measurements simultaneously at several neighborhood sizes, looking for agreement between the results obtained in those neighborhood sizes that lie just above the size at which random fluctuations appear. This technique can be applied to region finding or to boundary finding, and an example of the results is given in figure 5. The dot density here is not known *a priori*.


This issue is of considerable technical interest, but it is important not to lose sight of the underlying computational problem, which is what kind of boundary is to be found, and why? The techniques of O'Callaghan (1974) for example are designed to find boundaries in dot patterns so accurately that their positions are determined up to the decision about which dots it passes through. The justification for this type of study is that humans can assign boundaries this accurately, but the difficulty lies in formulating a reasonable definition of what the boundary is.

This problem is a deep one, touching the heart of the question of what early vision is for. I shall return to it later in this essay, but it is perhaps worth remarking here that there seems to be a clear need for being able to do early visual processing roughly and fast as well as more slowly and accurately, which means having ways of handling rough descriptions of regions — ways of characterizing their approximate extent and shape — before characterizing their precise boundaries. Figure 6 contains one example of a region whose rough extent is clear, but whose exact boundary is not.


The motivation for wanting this is that rough descriptions are very useful during the early stages of building a shape description for recognition (Marr & Nishihara 1977). For example a man often appears as a roughly vertical rectangle in an image, and this information is useful because it eliminates many other shapes from consideration quite early. Campbell (1977) has suggested that the extraction of rough descriptions from an image may depend on the ability to examine its lower spatial frequencies. Even if this is one of the available mechanisms it is unlikely to be the only one, because sparse line drawings can raise the same problems while having almost no power in their low frequencies. It may be that some notion of rough grouping applied to low resolution place-tokens set up by pieces of contour in the image provides a useful approach to this



5. Finding a boundary from dot (or place-token) density changes. Once a rough assignment of boundary points has been made (5a), local line-fitting (5b) and grouping (5c & d) techniques can recover a rough specification of the boundary quite easily.



6. An example of a region whose rough boundary is clear, but whose exact boundary is not.
(Drawing by K. Prendergast, 1977).



7. An image (7a), the spatial components of its primal sketch (7b), and a reconstruction of the image from the primal sketch (7c). This shows that the our current primal sketch programs lose little of the information in an image.

problem.

1.6: Lightness

Ever since Ernst Mach noticed the bands named after him, there has been considerable interest in the problem of computing perceived brightness. Of especial interest is the recent work of Land & McCann (1971) on the retinex theory (see also Horn 1974), which is concerned with the quantity they call lightness; and that of Colas-Baudelaire (1973) on the computation of perceived brightness. Lightness is an approximation to reflectance that is obtained by filtering out slow intensity changes, the underlying idea being that these are usually due to the illuminant, not to changes in reflectance. The problem with this idea is of course that some slow changes in intensity are perceptually important (see Horn 1977 for an analysis of shape from shading). The linear filter model of Colas-Baudelaire performs well on images in which there are no sharp changes in intensity, but the author found it difficult to extend his model to the more general case. The recent finding of Gilchrist (1977), that perceived depth influences perceived brightness, suggests that some aspects of the problem occur quite late -- in our terms, at the level of the $2\frac{1}{2}$ -D sketch (see below).

Our own work on the brightness problem is probably not relevant to the perception of brightness, but it is interesting as a demonstration that the primal sketch loses very little information. Woodham & Marr (unpublished program) have written a program that inverts the primal sketch, so that its output is an intensity array. The basic idea is to scan outwards from edges, assigning a constant brightness to points along the scan lines, and arresting the scan when it encounters another edge. Figure 7 exhibits the results of running this program, showing the original image (7a), the primal sketch (7b), and the reconstructed intensity array (7c).

2: Process-oriented theories

2.0: Introduction

I said earlier that, especially at the earlier stages of visual information processing, the representations and processes are determined more by what it is possible to compute from an image than by what is desirable. Examples are the problems associated with structure from motion, stereopsis, texture gradients, and shading.

2.1: Structure from motion

Given a sequence of views of objects in motion, the human visual system is capable of interpreting the changing views in terms of the shapes of the viewed objects, and their motion in three-dimensional space. Even if each successive view is unrecognizable, the human observer


easily perceives these views in terms of moving objects (Wallach & O'Donnell 1953). To answer the question of how a succession of images yields an interpretation in terms of three-dimensional structure in motion, Ullman (1977) divided the problem into two parts: (1) finding a correspondence between elements in successive views; and (2) determining the three-dimensional structures and their motion from the way corresponding elements move between views.

An important preliminary question about the correspondence problem concerns the level at which it takes place. Is it primarily a low-level relation, established between small and simple parts of the scenes and largely independent of higher-level knowledge and three-dimensional interpretation? Or do higher level influences, like the interpretation of the whole of a shape from one frame, play an important part in determining the correspondence?

Ullman has assembled a considerable amount of evidence that the former view is correct. For example, figure 8 shows two successive frames, one denoted with full lines and the other with dotted lines. If the whole pattern were being analyzed from one frame, the shape of the wheel extracted, and used to match the elements in the next frame, the observer presented with these frames in rapid succession should perceive them as a whole wheel rotating. Notice however that the inner and outer parts of the wheel have their closest neighbors in one direction, whereas the center parts have theirs in the other; because of this, if the matching were done early and locally, the observer should see the center part rotating one way, and the inner and outer rings rotating the other (as shown with arrows in figure 8). When appropriately timed, this is in fact what happens.

Another line of evidence is the following. The most important factor in finding a correspondence between elements is the distance the element moves from one view to the next. But is this distance an objective two-dimensional measurement or an interpreted movement in three-dimensional space? There is some confusion in the literature about this point, since many studies have assumed that correspondence strength is linked to the smoothness of apparent motion (Kolers 1972), and this is apparently more closely related to three- than to two-dimensional distances. Ullman (1977) has however shown that this assumption is false, and that it is the two-dimensional distance alone that determines the correspondence.

The second part of the problem is to determine the three-dimensional structure once the correspondence between successive views has been established. Unless this problem is constrained in some way, it cannot be solved, so one has to search for reasonable assumptions on which to base the design of one's algorithms. (This state of affairs is a common one in the theory of visual processes, as we shall see when we discuss the problems of stereopsis, and shape from contour). Ullman suggested basing the interpretation on the following assumptions: (1) any two-dimensional transformation that has a unique interpretation as a rigid body moving in space should be interpreted as such an object in motion, and (2) that the imaging process is locally an orthogonal projection. He then showed that under orthogonal projection, three-dimensional shape



8. Evidence that the correspondence problem for apparent motion involves matching operations that act at a low-level. Frame 1 is shown with full lines and frame 2 with dotted lines. Instead of seeing a single wheel rotating, when appropriately timed the wheel splits, the outer and inner rings rotating one way, and the center rotating the other, as indicated by the arrows. This suggests that matching is carried out on elemental line segments, and is governed primarily by proximity. (Adapted from Ullman 1977).

and motion may be recovered from as little as three views each showing the image of the same five points, no four of which are coplanar. This result leads to algorithms capable of recovering shape and motion from scenes containing arbitrary objects in motion. The final question is whether the algorithms that humans employ to recover shape and motion rely on these same two assumptions, and this question is currently under investigation. The important point here is that for more human-like algorithms, the number of views can be traded off against the accuracy of the computation, decreasing the emphasis on the particular number "three".

2.2: Stereopsis

Ever since Julesz (1971) made the first random-dot stereogram, it has been clear that at least to a first approximation stereo vision can be regarded as a modular component of the human visual system. Marr (1974) and Marr & Poggio (1976) formulated the computational theory of the stereo matching problem in the following way:




(R1) *Uniqueness.* Each item from each image may be assigned at most one disparity value. This condition rests on the premise that the items to be matched correspond to physical marks on a surface, and so can be in only one place at a time.

(R2) *Continuity.* Disparity varies smoothly almost everywhere. This condition is a consequence of the cohesiveness of matter, and it states that only a relatively small fraction of the area of an image is composed of boundaries.

By representing these constraints geometrically, Marr & Poggio (1976) embodied them in a cooperative algorithm. In figure 9, Lx and Rx represent the positions of descriptive elements from the left and right views, and the horizontal and vertical lines indicate the range of disparity values that can be assigned to left-eye and right-eye elements. The uniqueness condition then corresponds to the assertion that only one disparity value may be "on" along each horizontal or vertical line. The continuity condition states that we seek solutions that tend to spread along the dotted diagonals, which are lines of constant disparity, and between adjacent diagonals. Figure 9b shows how this geometry appears at each intersection point. Figure 9c gives the corresponding local geometry when the images are two-dimensional rather than one. and R2 for the case of a one-dimensional image, and it also represents the structure of a network for implementing the algorithm described by equation 1. Solid lines represent "inhibitory" interactions, and dotted lines represent "excitatory" ones. 9b gives the local structure at each node of the network 9a. This algorithm may be extended to two-dimensional images, in which case each node in the corresponding network has the local structure shown in 9c. Such a network was used to solve the stereograms exhibited in figures 10 and 11. (Marr & Poggio 1976 figure 2).

It can be shown (Marr, Poggio & Palm 1977) that, if a network is created with the positive and negative connections shown in figure 9c, states of such a network that satisfy the constraints on the computation are stable, and that given suitable inputs, the network will converge









9. Figure 9a shows the explicit structure of the two rules R_1 and R_2 for the case of a one-dimensional image, and it also represents the structure of a network for implementing the algorithm described by equation 1. Solid lines represent "inhibitory" interactions, and dotted lines represent "excitatory" ones. 9b gives the local structure at each node of the network 9a. This algorithm may be extended to two-dimensional images, in which case each node in the corresponding network has the local structure shown in 9c. Such a network was used to solve the stereograms exhibited in figures 10 and 11. (Marr & Poggio 1976 figure 2).



10. This and the following figure show the results of applying the algorithm defined by equation (1) to two random-dot stereograms. The initial state of the network C_{xyd} is defined by the input such that a node takes the value 1 if it occurs at the intersection of a 1 in the left and right eyes (see figure 9), and it has value 0 otherwise. The network iterates on this initial state, and the parameters used here, as suggested by the combinatorial analysis, were $\theta = 3.0$, $\epsilon = 2.0$ and $M = 5$, where θ is the threshold and M is the diameter of the "excitatory" neighborhood illustrated in figure 9c. The stereograms themselves are labelled LEFT and RIGHT, the initial state of the network as 0, and the state after n iterations is marked as such. To understand how the figures represent states of the network, imagine looking at it from above. The different disparity layers in the network lie in parallel planes spread out horizontally, so that the viewer is looking down through them. In each plane, some nodes are on and some are off. Each of the seven layers in the network has been assigned a different gray level, so that a node that is switched on in the top layer (corresponding to a disparity of +3 pixels) contributes a dark point to the image, and one that is switched on in the lowest layer (disparity = -3) contributes a lighter point. Initially (iteration 0) the network is disorganized, but in the final state, stable order has been achieved (iteration 14), and the inverted wedding-cake structure has been found. The density of this stereogram is 50%. (Marr & Poggio 1976 figure 9).

LEFT

RIGHT



ii. The algorithm of equation 1, with parameter values given in the legend to figure 10, is capable of solving random-dot stereograms with densities from 50% down to less than 10%. For this and smaller densities, the algorithm converges increasingly slowly. If a simple homeostatic mechanism is allowed to control the threshold θ as a function of the average activity (number of "on" cells) at each iteration, the algorithm can solve stereograms whose density is very low. In this example, the density is 5% and the central square has a disparity of +2 relative to the background. The algorithm "fills in" those areas where no dots are present, but it takes several more iterations to arrive near the solution than in cases where the density is 50%. When we look at a sparse stereogram, we perceive the shapes in it as cleaner than those found by the algorithm. This seems to be due to subjective contours that arise between dots that lie on shape boundaries. (Marr and Poggio 1976 figure 4).

LEFT

RIGHT

0

1

2

3

4

5

7

10

13

to these stable states for a wide variety of the control parameters. Thus one can think of the network as defining an algorithm that operates on many input elements to produce a global organization *via* local but highly interactive constraints. Formally, the algorithm reads:

$$C_{xyd}^{(n-1)} = u \left\{ \sum_{x'y'd' \in S(xyd)} C_{xyd}^{(n)} - \epsilon \sum_{x'y'd' \in O(xyd)} C_{x'y'd'}^{(n)} + C_{xyd}^{(o)} \right\} \quad (1)$$

where $u(z) = 0$ if $z < 0$, and $u(z) = 1$ otherwise; S and O are the circular and thick line neighborhoods of the cell C_{xyd} in figure 9c. This is an example of a "cooperative" algorithm (Marr & Poggio 1977a), and it exhibits typical non-linear cooperative phenomena like hysteresis, filling-in, and disorder-order transitions. Figures 10 and 11 illustrate two applications of the algorithm to random-dot stereograms.

There are a number of findings that cast doubt on the relevance of this algorithm to the question of how human stereo vision works. The most important of these findings are (a) the apparently crucial role played by eye-movements in human stereo vision (see especially Richards 1977); (b) our ability to tolerate up to 15% expansion of one image (Julesz 1971 figure 2.8.8); (c) our ability to tolerate the severe defocussing of one image (Julesz 1971 figure 3.10.3); (d) evidence that stereo detectors are organized into "three pools" (convergent, zero disparity, and divergent) and that this organization is important for stereo vision (Richards 1971); and (e) our ability to perceive depth in rivalrous stereograms (Mayhew & Frisby 1976). These difficulties led Marr & Poggio (1977b) to formulate a second stereo algorithm, designed specifically as a model for human stereopsis.

Our first stereo theory was inspired by Julesz's belief that stereoscopic fusion is a cooperative process -- a belief based primarily on the observation that it exhibits hysteresis. The main problem with the cooperative algorithm is that it apparently works too well in some ways (it performs better than humans do when eye-movements are eliminated), and not well enough in others (humans see depth in rivalrous stereograms). Our ability to fuse two images when one is blurred, the rivalrous stereogram results of Mayhew & Frisby (1976), and the recent results of Julesz & Miller (1975) on the existence of independent spatial-frequency-tuned channels in binocular fusion, suggest that several copies of the image, obtained by successively coarser filtering, are used during fusion, perhaps helping one another in a way similar to that in which local regions help each other in our cooperative algorithm.

The second idea was a notion that originated with Marr & Nishihara (1977) and about which I shall have more to say later, which is that one of the things early visual processing does is to construct a "depth map" of the surfaces round a viewer. In this map, each direction away from the viewer is associated with a distance (or some function of distance) and a surface orientation. We have christened the resulting datastructure the $2\frac{1}{2}$ -D sketch.

The important point here is that the $2\frac{1}{2}$ -D sketch is in some sense a memory. This

provided the key idea: Suppose that the hysteresis Julesz observed is not due to a cooperative process at all, but is in fact the result of using a memory buffer in which to store the depth map of the image as it is discovered. Then, the fusion process itself need not be cooperative, and in fact it would not even be necessary for the whole image ever to be fused everywhere provided that a depth map of the viewed surface were built and maintained in this intermediate memory. This idea leads to the following theory. (1) Each image is convolved with bar-shaped masks of various sizes, and matching takes place between peak mask values for disparities up to about twice the panel-width of the mask (see Felton, Richards & Smith 1972), for pairs of masks of the same size and polarity. (2) Wide masks can control vergence movements, thus causing small masks to come into correspondence. (3) When a correspondence is achieved, it is held and written down somewhere (e.g. in the $2\frac{1}{2}$ -D sketch). (4) There is a backwards relation between the memory and the masks, perhaps simply through the control of eye-movements, that allows one to fuse any piece of a surface easily once its depth map has been established in the memory.

This theory leads to many experimental predictions, which are currently being tested.

8: Intermediate processing problems

3.0 Introduction

We have discussed the types of information that need to be represented early in the processing of visual information, and we have examined the computational structure of some of the processes that can derive and maintain this information. We turn now to the question of what all this information is to be used for.

3.1 Difficulties with the idea of image segmentation

The current approach to machine vision assumes that the next step in visual processing consists of a process called *segmentation*, whose purpose is to divide the image into regions that are meaningful either in terms of physical objects or for the purpose at hand. Despite considerable efforts over a long period, the theory and practise of segmentation remain primitive, and once again I believe that the main reason lies in the failure to formulate precisely the goals of this stage of the processing. What for example is an object? Is a head one? Is it still one if it is attached to a body? What about a man on horseback?

These questions point to some of the difficulties one has when trying to formulate what should be recovered as a region from early visual processing. Furthermore, however one chooses to answer them, it is usually still impossible to recover the desired regions using only local grouping techniques acting on a representation like the primal sketch. Most images are too

complex, and even the simplest images cannot often be segmented entirely at that level (e.g. Marr 1976 figure 13).

Something additional is clearly needed, and one approach to the dilemma has been to invoke specialized knowledge about the nature of the scenes being viewed to aid segmentation of the image into regions that correspond roughly to the objects expected in the scene. Tenenbaum & Barrow (1978), for example, applied knowledge about several different types of scene to the segmentation of images of landscapes, an office, a room, and a compressor. Freuder (1974) used a similar approach to identify a hammer in a simple scene. If this approach were correct, it would mean that a central problem for vision is arranging for the right piece of specialized knowledge to be made available at the appropriate time during segmentation. Freuder's work, for example, was almost entirely devoted to the design of a hierarchical control system that made this possible. More recently, the constraint relaxation technique of Rosenfeld, Hummel & Zucker (1978) has attracted considerable attention for just this reason, that it appears to offer a technique whereby constraints drawn from disparate sources may be applied to the segmentation problem whilst incurring only minimal penalties in control. It is however difficult to analyze such algorithms rigorously even in very clearly defined situations (see e.g. Marr, Poggio & Palm 1977), and in the naturally more diffuse circumstances that surround the segmentation problem, it may often be impossible.

3.2: Reformulating the problem

The basic problem seems to be how to formulate precisely the next stage of visual processing. Given a representation like the primal sketch, and the many possible boundary-defining processes that are naturally associated with it, which boundaries should one attend to and why? The segmentation approach fails because objects and desirable regions are not visually primitive constructions, and hence cannot be recovered reliably from the primal sketch or similar representation without additional specialized knowledge. If we are to succeed, we must discover precisely what information it is that needs to be made explicit at this stage, what, if any, additional knowledge it is appropriate to apply, and we must design a representation that matches these requirements.

In order to search for clues to a suitable representation, let us return to the physics of the situation. The primal sketch represents intensity changes and the local two-dimensional geometry of an image. The principle factors that determine these are (1) the illuminant, (2) surface reflectance, (3) the shape of the visible surface, and (4) the vantage point. The first two factors raise the difficult problems of color and brightness, and I shall not discuss them further. The third and fourth factors are independent of the first two (whether two shapes are the same does not depend upon their colors or on the lighting), and so may be treated separately.

I shall argue that, since most early visual processes extract information about the

visible surface, it is these surfaces, their shape and disposition relative to the viewer, that need to be made explicit at this point in the processing. Furthermore, because surfaces exist in three-dimensional space, this imposes constraints on them that are general, and not confined to particular objects. It is these constraints that constitute the *a priori* knowledge that it is appropriate to bring to bear next.

One example of the exploitation of fairly general constraints was the work of Waltz (1975), who formulated the constraints that apply to images of polyhedra. The representation on which that work was based was line drawings, but these are not suitable for our needs here, because part of the task we wish to carry out is the discovery of physical edges that are only weakly present or even absent in the primal sketch. The approach of Mackworth (1979) was closer to what we want, since it involved a primitive way of representing surfaces.

3.3: General classification of shape representations

Part of our task in formulating the problem of intermediate vision is therefore the examination of ways of representing and reasoning about surfaces. We therefore start our enquiry by discussing the general nature of shape representations. What kinds are there, and how may one decide among them? Although it is difficult to formulate a completely general classification of shape representations, Marr & Nishihara (1977) attempted to set out the basic design choices that have to be made when a representation is formulated. They concluded that there are three characteristics of a shape representation that are largely responsible for determining the information that it makes explicit. The first is the type of coordinate system it uses, whether it is defined relative to the viewer or to the object being viewed; the second characteristic concerns the nature of the *shape primitives* used by the representation, that is, the elements whose positions the coordinate system is used to define. Are they two- or three-dimensional, in what sizes do they come, and how detailed are they? And the third is concerned with the organization a representation imposes on the information in a description, for example is the description modular or does it have little internal structure? We have two sources of information that can help us to formulate the important issues in intermediate visual information processing, firstly the computational problems that arise, and secondly, psychophysics.

3.4: Some observations from psychophysics

Vision provides several sources of information about shape. The most direct are stereo and motion, but texture gradients in a single image are nearly as effective, and the theatrical techniques of facial make-up rely on the sensitivity of perceived shape to shading. It often happens that some parts of a scene are open to inspection by some of these techniques, and other parts by others. Yet different as the techniques are, they have two important characteristics in common. They rely on information from the image rather than on *a priori* knowledge about the

shapes of the viewed objects; and the information they specify concerns the depth or surface orientation at arbitrary points in an image, rather than the depth or orientation associated with particular objects.

If one views a stereo pair of a complex surface, like a crumpled newspaper or the "leaves" cube of Ittelson (1960), one can easily state the surface orientation of any piece of the surface, and whether one piece is nearer to or further from the viewer than its neighbors. Nevertheless one's memory for the shape of the surface is poor, despite the vividness of its surface orientation during perception. Furthermore, if the surface contains elements nearly parallel to the line of sight, their apparent surface orientation when viewed monocularly can differ from the apparent surface orientation when viewed binocularly.

From these observations, one can perhaps draw some simple inferences.

(a) There is at least one internal representation of the depth, or surface orientation, or both, associated with each surface point in a scene.

(b) Because surface orientation can be associated with unfamiliar shapes, its representation probably precedes the decomposition of the scene into objects. (This point is particularly relevant to our discussion of intermediate visual information processing.)

(c) Because the apparent orientation of a surface element can change, depending on whether it is viewed binocularly or monocularly, the representation of surface orientation is probably driven almost entirely by perceptual processes, and is influenced only slightly by specific knowledge of what the surface orientation actually is. Our ability to "perceive" the surface much better than we can "memorize" it may also be connected with this point.

In addition, it seems likely that the different sources of information can influence the same representation of surface orientation.

3.5: The computational problem

In order to make the most efficient use of these different and often complementary sources of information, they need to be combined in some way. The computational question is, how best to do this? The natural answer is to seek some representation of the visual scene that makes explicit just the information these processes can deliver.

Fortunately, the physical interpretation of the representation we seek is clear. All these processes deliver information about the depth or surface orientation associated with surfaces in an image, and these are well-defined physical quantities. We therefore seek a way of making this information explicit, of maintaining it in a consistent state, and perhaps also of incorporating into the representation any physical constraints that hold for the values that depth and surface orientation take over the kinds of surface that occur in the real world. Table I lists the type of information that the different early processes can extract from images. The interesting point here is that although processes like stereo and motion are in principle capable of delivering depth

Table 1

The form in which various early visual processes deliver information about the changes in a scene.

r = depth

δr = small, local changes in depth

Δr = large changes in depth

s = local surface orientation

<u>Information source</u>	<u>Natural parameter</u>
Stereo	Disparity, hence especially δr and Δr
Motion	r , hence δr , Δr
Shading	s
Texture gradients	s
Perspective cues	s
Occlusion	Δr



information directly, they are in practise more likely to deliver information about local changes in depth, for example by measuring local changes in disparity. Texture gradients and shading provide more direct information about surface orientation. In addition, occlusion and brightness and size clues can deliver information about discontinuities in depth. (It is for example amazing how clear an impression of depth can be obtained from a monocular image containing bright or dim rectangles of different sizes against a dark background). The main function of the representation we seek is therefore not only to make explicit information about depth, local surface orientation, and discontinuities in these quantities, but also to create and maintain a global representation of depth that is consistent with the local cues that these sources provide. We call such a representation the $2\frac{1}{2}$ -D sketch, and the next section describes a particular candidate for it.

3.6: A possible form for the $2\frac{1}{2}$ -D sketch

The example I give for the $2\frac{1}{2}$ -D sketch is a viewer-centered representation, which uses surface primitives of one (small) size. It includes a representation of contours of surface discontinuity, and it has enough internal computational structure to maintain its descriptions of depth, surface orientation and surface discontinuity in a consistent state. The representation itself has no additional internal structure.

Depth may be represented by a scalar quantity r , the distance from the viewer of a point on a surface. Surface discontinuities may be represented by oriented line elements. Surface orientation may be represented by a unit vector (x, y, z) in three-dimensional space. Following those who have used gradient space (Huffman 1971, Horn 1977) we can rewrite this as $(p, q, 1)$, which can be represented as a vector (p, q) in two-dimensional space. In other words, surface orientation may be represented by covering an image with needles. The length of each needle defines the dip of the surface at that point, so that zero length corresponds to a surface that is perpendicular to the vector from the viewer to that point, and the length increases as the surface tilts away from the viewer. The orientation of the needle defines the direction of the surface's dip. Figure 12 illustrates this representation.

In principle, the relation between depth and surface orientation is straightforward — one is simply the integral of the other, taken over regions bounded by surface discontinuities. It is therefore possible to devise a representation with intrinsic computational facilities that can maintain the two variables, of depth and surface orientation, in a consistent state. But note that, in any such scheme, *surface discontinuities* acquire a special status (as curves across which integration stops). Furthermore, if the representation is an active one, maintaining consistency through largely local operations, curves that mark surface discontinuities (e.g. contours that arise from occluding contours in the image) must be "filled in" completely, so that at no point along an object boundary can the integration leak across it. It is interesting that subjective contours have this property, and that they are closely related to subjective changes in brightness (cf section 1.6) that are often



12. The $2\frac{1}{2}$ -D sketch represents depth, contours of surface discontinuity, and the orientation of visible surfaces. A convenient representation of surface orientation is described in the text and illustrated here. The orientation of the needles is determined by the projection of the surface normal on the image plane, and the length of the needles represents the dip out of that plane (12a). A typical $2\frac{1}{2}$ -D sketch appears in 12b, although depth information is not represented in the figure. (Marr & Nishihara 1977 figure 2).

associated with changes in perceived depth. If the human visual processor contains a representation that resembles the $2\frac{1}{2}$ -D sketch, it would therefore be interesting to ask whether subjective contours occur within it. (See Ullman (1978) for an analysis of the shape of curved subjective contours).

In summary, my argument is that the $2\frac{1}{2}$ -D sketch is useful because it makes explicit information about the image in a form that is closely matched to what early visual processes can deliver. We can formulate the goals of intermediate visual processing as being primarily the construction of this representation, discovering for example what are the surface orientations in a scene, which of the contours in the primal sketch correspond to surface discontinuities and should therefore be represented in the $2\frac{1}{2}$ -D sketch, and which contours are missing in the primal sketch and need to be inserted into the $2\frac{1}{2}$ -D sketch in order to bring it into a state that is consistent with the structure of three-dimensional space. This formulation avoids the difficulties associated with the terms "region" and "object", and allows one to ask precise questions about the computational structure of the $2\frac{1}{2}$ -D sketch and of processes to create and maintain it. We are currently much occupied with these problems.

4: Later processing problems

4.0: Introduction

The $2\frac{1}{2}$ -D sketch is a poor representation for the purposes of recognition because it is unstable (in the sense of Marr & Nishihara 1977), it depends on the vantage point, and it fails to make explicit pieces of a shape (like an arm) that are larger than the primitive size. Except for the simplest of purposes, it is an inadequate vehicle for a visual system to convey information about shape to other processes, and so I turn now to representations that are more suitable for recognition tasks.

If one were to design a shape representation to suit the problems of recognition, one would naturally base it on an object-centered coordinate system. In addition, one would have to include shape primitives of many different sizes, so as to be able to make explicit shape characteristics that can range from a wart to an elephant. Marr & Nishihara (1977) discuss these questions in detail, and I shall not repeat their observations here. The deepest issues are those raised by having to define an object-based coordinate system. Since they are central to the problem of defining representations for use in later processing of visual information, I shall spend the remainder of the essay discussing this topic.

4.1: Nature of an object-centered coordinate system


Marr & Nishihara (1977) pointed out that there are two types of object-centered coordinate system that one might attempt to define precisely. One refers all locations on an object to a single coordinate frame that embraces the entire object, and the other distributes the coordinate system, making it local to each articulated component or individual shape characteristic. Marr & Nishihara concluded that the second of these schemes is the more desirable, and they gave as an example the representation illustrated in figure 13. But with a representation of this kind, the most difficult questions begin after its internal structure has been defined. How can one define canonically the coordinate scheme for an arbitrary shape, and even more difficult, how can such a thing be found from an image before a description of the viewed shape has been computed? Some kind of answers to these questions must be found if the representation is to be used for recognition.

4.2: Shapes having natural coordinate systems


If the coordinate system used for a given shape is to be canonical, its definition must take advantage of any salient geometrical characteristics that the shape possesses. For example, if a shape has natural axes, distinguished by length or by symmetry, then they should be used: The coordinate system for a sausage should take advantage of its major axis, and for a face, of its axis of symmetry.

Highly symmetrical objects, like a sphere, square, or circular disc, will inevitably lead to ambiguities in the choice of coordinate systems. For a shape as regular as a sphere this poses no great problem, because its description in all reasonable systems is the same. One can even allow other factors, like the direction of motion or of spin, to influence the choice of coordinate frame. For other shapes, the existence of more than one possible choice probably means that one has to represent the object in several ways. This is acceptable provided that the number of ways is small. For example, there are four possible axes on which one might wish to base the coordinate system for representing a door, the midlines along its length, its width, its thickness, and to represent how the door opens, the axis of its hinges. For a typewriter, there are two choices at the top level; an axis parallel to its width, because that is usually its largest dimension, and the axis about which a typewriter is roughly symmetrical.

In general, if an axis can be distinguished in a shape, it can be used as the basis for a local coordinate system. One approach to the problem of defining object-centered coordinate systems is therefore to examine the class of shapes having an axis as an integral part of their structure. One such is the class of *generalized cones*. (A generalized cone is the surface swept out by moving a cross section of constant shape but smoothly varying size along an axis, as in figure 14). Binford (1971) drew attention to this class of surfaces, suggesting that it might provide a




13. This diagram illustrates the organization of shape information in a 3-D model description. Each box corresponds to a 3-D model; with its model axis on the left side of the box and the arrangement of its component axes are shown on the right side. In addition some component axes have 3-D models associated with them and this is indicated by the way the boxes overlap. The relative arrangement of each model's component axes, however, is shown improperly since it should be in an object-centred system rather than the viewer-centred projection used here. This example shows a coarse overall description of a human shape along with an elaboration of one of its components (the arm). The important characteristics of this type of organization are: (i) each 3-D model is a self-contained unit of shape information and has a limited complexity, (ii) information appears in shape contexts appropriate for recognition (the disposition of a finger is most stable when specified relative to the hand that contains it), and (iii) the representation can be used flexibly (components can be elaborated according to the needs of the moment or the time available, and a 3-D model description of a component is easily added to a description of the whole shape). The major limitation imposed on the representation by this form of organization is on its scope, since it will only be useful for shapes for which the decomposition into 3-D models is well defined. (Marr & Nishihara 1977 figure 3).



14. The definition of a generalized cone. In this article, a generalized cone is the surface generated by moving a smooth cross-section ρ along a straight axis Λ . The cross-section may vary smoothly in size (as prescribed by the function $\Lambda(z)$), but its shape remains constant. The eccentricity of the cone is the angle ψ between its axis and a plane containing a cross-section. (Figure 5 of Marr 1977).

15. These pipecleaner figures illustrate the point that a shape representation does not have to reproduce a shape's surface in order to describe it adequately for recognition; as we see here, animal shapes can be portrayed quite effectively by the arrangement and relative sizes of a small number of sticks. The simplicity of these descriptions is due to the correspondence between the sticks shown here and natural or canonical axes of the shapes described. To be useful for recognition, a shape representation must be based on characteristics that are uniquely defined by the shape and which can be derived reliably from images of it. (Marr & Nishihara 1977 figure 1).




convenient way of describing three-dimensional surfaces for the purposes of computer vision. I regard it as an important class not because the shapes themselves are easily describable, but because the presence of an axis allows one to define a canonical local coordinate system. Fortunately many objects, especially those whose shape was achieved by growth, are described quite naturally in terms of one or more generalized cones. The animal shapes in figure 15 provide some examples -- the individual sticks are simply axes of generalized cones that approximate the shapes of parts of these animals. Many artifacts can also be described in this way, like a car (a small box sitting atop and in the middle of a longer one), and a building (a box with a vertical axis).

It is important to remember that there exist surfaces that cannot conveniently be approximated by generalized cones, for example a cake that has been cut at its intersection with some arbitrary plane, or the surface formed by a crumpled newspaper. Cases like the cake can be dealt with by introducing suitable surface primitives that describe the plane of the cut, but the crumpled newspaper poses apparently intractable problems.

4.3: Finding the natural coordinate system from an image

Even if a shape possesses a canonical coordinate system, one is still faced with the problem of finding it from an image. Blum (1973), Agin (1972) and Nevatia (1974) have addressed problems that are related to this question. Blum's sym-axis theory is an interesting one, because he specifies precisely what it is that is computed from a two-dimensional outline. Unfortunately, it is not clear that what this theory computes is in fact useful for shape recognition (see e.g. figure 16), and when applied to a three-dimensional shape, the sym-axis is in general a two-dimensional sheet, so it cannot easily be used to define an object-centered coordinate system. Agin's and Nevatia's work, on the other hand, concerns the analysis of a depth map. This is an important problem, and it would be interesting to see a careful analysis of the conditions under which their techniques will succeed.

My own interest in the problem grew from the 3-D representation theory of Marr & Nishihara (1977), in particular from the question of how to interpret the outlines of objects as seen in a two-dimensional image. The rest of this essay summarizes a recent article by Marr (1977). The starting point for this work was the observation that when one looks at the silhouettes in Picasso's work "Rites of Spring" (figure 17), one perceives them in terms of very particular three-dimensional shapes, some familiar, some less so. This is quite remarkable, because the silhouettes could in theory have been generated by an infinite variety of shapes which, from other viewpoints, have no discernable similarities to the shapes we perceive. One can perhaps attribute part of the phenomenon to a familiarity with the depicted shapes; but not all of it, because one can use the medium of a silhouette to convey a new shape, and because even with considerable effort it is difficult to imagine the more bizarre three-dimensional surfaces that could have given rise to the same silhouettes. The paradox is, that the bounding contours in figure 17 apparently tell us more








16. Blum's (1979) grassfire technique for recovering an axis from a silhouette is undesirably sensitive to small perturbations in the contour. 16a shows the Blum transform of a rectangle, and 16b, of a rectangle with a notch. (Redrawn from Agin 1972).

17. "Rites of spring" by P. Picasso. We immediately interpret the silhouettes in terms of particular 3-D surfaces, despite the paucity of information in the image. In order to do this, we must be bringing additional assumptions and constraints to bear on the analysis of these contours' shapes. Marr (1977) enquired about the nature of this *a priori* information.

~~Picasso~~
13.11.57



18. From viewpoint V , the three-dimensional surface Σ forms the silhouette S_V in the image via the imaging process ι . The boundary of S_V , obtained by the boundary operator ∂ is denoted by C_V and we call it the contour of Σ . The set of points on Σ that ι maps onto C_V we call the contour generator of C_V , and it is denoted by Γ_V . The map from Σ to Γ_V induced by ∂ is denoted by δ . (Figure 2 of Marr 1977).



than they should about the shape of the dark figures. For example, neighboring points on such a contour could in general arise from widely separated points on the original surface, but our perceptual interpretation usually ignores this possibility.

The first observation to be made here is that the occluding contours that bound these silhouettes are contours of surface discontinuity, that is precisely the contours with which the $2\frac{1}{2}$ -D sketch is concerned. Secondly, because we can interpret them as three-dimensional shapes, then implicit in the way we interpret them must lie some *a priori* assumptions that allow us to infer a shape from an outline. If a surface violates these assumptions, our analysis will be wrong, in the sense that the shape we assign to the contours will differ from the shape that actually caused them. An everyday example of this phenomenon is the shadowgraph, where the appropriate arrangement of one's hands can, to the surprise and delight of a child, produce the shadow of an apparently quite different shape, like a duck or a rabbit.

What assumptions is it reasonable to suppose that we make? In order to explain them, I need to define the four structures that appear in figure 18. These are (1) some three dimensional surface Σ ; (2) its image or silhouette S_V as seen from a viewpoint V ; (3) the bounding contour C_V of S_V ; and (4) the set of points on the surface Σ , that project onto the contour C_V . We shall call this last set the *contour generator* of C_V , and we shall denote it by Γ_V .

If one is presented with a contour in an image, without any knowledge of the surface or perspective that caused it, there is very little information on which one can base one's analysis. The only obvious feature available is the distinction between convex and concave pieces of contour — that is, the presence of inflection points. In order that inflection points be "reliable", one needs to make some assumptions about the way the contour was generated, and I chose the following restrictions:

R1: The surface Σ is smooth.


R2: Each point on the contour generator Γ_V projects to a different point on the contour C_V .

R3: Nearby points on the contour C_V arise from nearby points on the contour generator Γ_V .

R4: The contour generator Γ_V of C_V is planar.

The first restriction is only a technical one. The second and third say that each point on the contour in the image comes from one point on the surface (which is an assumption that facilitates the analysis but is not of fundamental importance), and that where the surface looks continuous in the image, it really is continuous in three dimensions. The fourth condition, together with the constraint that the imaging process be an orthogonal projection, is simply a necessary and sufficient condition that the difference between convex and concave contour segments reflects properties of the surface, rather than characteristics of the imaging process.

It turns out that the following theorem is true, and it is a result that I found very



19. A sketch of a generalized cone showing its silhouette (the circumscribing contour), and its fluting (the contours spanning its length). The radial extremities of a generalized cone are illustrated in Figure 20.

surprising.

Theorem. If $R1$ is true, and $R2 - R4$ hold for all distant viewing directions that lie in some plane, then the viewed surface is a generalized cone.

This means that if, for distant viewpoints whose viewing directions lie parallel to some plane, a surface's shape can successfully be inferred using only the convexities and concavities of its bounding contours in an image, then that surface is a generalized cone or is composed of several such cones. The interesting thing about this result is that it implicates generalized cones. We have already seen that the important thing about these cones is that an axis forms an integral part of their structure. But this is a feature of their three-dimensional organization, and ought in some sense to be independent of the issues raised by vision. What the theorem says is that there is a natural link between generalized cones and the imaging process itself. The combination of these two must mean, I think, that generalized cones will play an intimate role in the development of vision theory.

4.4: Interpreting the image of a single generalized cone


If we take this result at face value, we can now ask an obvious question. Let us assume that our data consist of contours of surface discontinuity in the image of a generalized cone, since without this assumption we can deduce nothing. How may such contours be interpreted? To specify a generalized cone, we have to specify its axis A , cross-section $\rho(\theta)$, and axial scaling function $\lambda(z)$ (figure 14); how can we discover them from an image?

The answer to this question is based on the notion of the *skeleton* of a generalized cone. The skeleton is not a difficult idea, since it is very like the set of lines a cartoonist draws to convey the shape of a curved object. It consists of three classes of contour: (a) the contours that occur in a generalized cone's silhouette; (b) the contours that arise from maxima and minima in a cone's axial scaling function (called the cone's *radial extremities*); and (c) contours that arise from maxima and minima in the cone's cross-section (its *fluting*). These categories are illustrated in figure 19.


The reason why the skeleton is a useful construct for recognition is that one can detect its presence in an image by the many relationships that exist among its parts. For example, radial extremities are all parallel to each other, and the silhouette and fluting have a kind of symmetry about the image of the cone's axis. It turns out that one can use these relationships to set up constraints on a set of contours such that, if those constraints

20. Methods based on the theory described in the text suffice to solve this image of a bucket. An axial symmetry is established by its sides about the bucket's axis (shown thickened), and a parallel relationship holds between components of its radial extremity. (Figure 14 of Marr 1977).


b.




a.




a.



b.



21. The two main types of joins between two generalized cones. 21a shows a side-to-end join, and 21b shows an end-to-end join. (Figure 14 of Marr 1977).



22. This figure illustrates the types of side-to-end join that can occur between two short generalized cones. In the first column, the left-hand cone is convex; in the center column it is concave, and in the third column, it is convex on one side of the join, and concave on the other. The other cone is convex in the top row, and concave in the other two. Segmentation depends upon finding the points P and Q , which are defined by theorem 7 of Marr (1977) and illustrated here for each case. (Figure 18 of Marr 1977).

are all satisfied by a unique interpretation of the contours in the image, one can be reasonably certain that a skeleton has been found, and hence that the contours can be interpreted as arising from a generalized cone whose axis is then determined. The practical importance of this result is illustrated in figure 20, where one can see that the image of the "sides" is symmetrical about the bucket's axis, and there is a clear parallel relationship between the image of the bucket's top, the corrugations in its side, and the visible part of its base (the bucket's radial extremities). These relations, of symmetries and parallelism, are preserved by an orthogonal projection. Hence provided that the contours are formed along a viewing direction that is not too close to the axis of the cone, these relations will still be present in the image. If the viewing direction lies so close to the cone's axis that its image is substantially foreshortened, these relationships will no longer be present, but it is part of the overall theory that such views have to be handled differently (Marr & Nishihara 1977).

4.5: Surfaces composed of two or more generalized cones

Real-life objects are often approximately composed of several different cones, joined together in various ways (see figure 19), and we therefore have to study ways of decomposing a multiple cone into its components – for example, a human body into arms, legs, torso and head. Marr (1977) analyzed the two types of join shown in figure 21, giving criteria that define segmentation points on the contour produced by two joined cones (theorems 7 and 8). Figure 22 exhibits the segmentation points P and Q for the case in which two short cones are joined side-to-end. P. Vatan has written a computer program that can carry out this segmentation, and an example of its operation is illustrated in figure 23. The legend to the figure describes the particular algorithm used.

4.6: Some comments on the limitations of this theory

The results of this theory are limited in their scope to a particular class of views and surfaces, but on the other hand, they use only a limited kind of visual information, little more than occluding contours that are formed in an image by rays that graze a smooth surface. Interestingly, these particular contours are unsuitable for use in stereopsis or structure-from-motion computations, because they are not formed from markings that define precise locations on the viewed surface. Creases and folds on a surface also give rise to contours in an image, and these have yet to be studied in detail. Information about shape from shading, texture, stereo or motion information has not yet been considered. By adding these other sources of information, I hope that a set of methods can eventually be assembled that together approach a comprehensive treatment of possible image configurations.

23. The occluding contours in an image can be used to locate the images of the natural axes of a shape composed of generalized cones (Marr 1977). The initial outline in (a) was obtained by applying local grouping processes to the *primal sketch* of the image of a toy donkey (Marr 1976). This outline was then smoothed and divided into convex and concave sections to get (b). Next, strong segmentation points, like the deep concavity circled in (c), are identified and a set of heuristic rules are used to connect them with other points on the contour to get the segmentation shown in (d). The component axes shown in (e) are then derived from these. The resulting segments are checked to see that they obey the rules for images of generalized cones. The boundaries must for example be symmetric about the axes, and in the case of side-to-end joins, the axis of the cone that is attached by its end must intersect the segmentation points that separate the two cones' contours. In this example, most of the symmetry relations have degenerated into parallelism. The thin lines in (f) indicate the position of the head, leg, and tail components along the torso axis, and the snout and ear components along the head axis. (This algorithm is due to P. Vatan).










TABLE 2

A framework for the derivation of shape information from images.



5: Discussion

I have tried to make two main points in this article. The first is methodological, namely that it is important to be very clear about the nature of the understanding we seek. The results we try to achieve should be precise ones, at the level of what I called computational theory, and one should try to choose problems that can confidently be attributed to a real aspect of vision, and not (for example) an artifact of the limitations of one's current vision program.

The second main point is that the critical issues for vision seem to me to revolve around the nature of the representations used, and the nature of the processes that create, maintain and read them. I have suggested an overall framework for visual information processing that consists of three principal representations, the primal sketch, the $2\frac{1}{2}$ -D sketch and the 3-D model representation, and it is summarized in table 2. As we study the processes capable of creating, maintaining and reading these representations, it is essential to make explicit their inherent limitations together with the assumptions that are implicit in their design. In addition, one should try hard to test experimentally any conclusions to which these studies lead, because it is foolish to ignore the clues and tests available from the disciplines of experimental psychology, neurophysiology and clinical neurology, even though it is often difficult to use this information fruitfully.

Acknowledgements: I thank the members of the M.I.T. Artificial Intelligence Laboratory vision group for making this survey possible, especially Keith Nishihara with whom much of the overall framework was developed, Dr. Tomaso Poggio and Prof. Whitman Richards for many stimulating discussions, and Karen Prendergast for preparing the illustrations. *Science* kindly gave permission for the reproduction of figures 9, 10 and 11, and the Royal Society for figures 1, 12, 13, 14, 15, 17, 18, 20, 21 and 22. This work was conducted at the Artificial Intelligence Laboratory, a Massachusetts Institute of Technology research program supported in part by the Advanced Research Projects Agency of the Department of Defense, and monitored by the Office of Naval Research under contract number N00014-75-C-0843.

6: References

- Agin, G. J. (1972) Representation and description of curved objects. *Stanford A.I. Memo 179.*

- Beck, J. (1974). *Surface color perception*. Ithaca, N. Y.: Cornell University Press.
- Binford, T. O. (1971) Visual perception by computer. Presented to the IEEE Conference on Systems and Control, Miami, December.
- Blum, H. (1973) Biological shape and visual science, (part I). *J. theor. Biol.*, 38, 205-287.
- Campbell, F. W. C. (1977) Sometimes a biologist has to make a noise like a mathematician. *NRP Bulletin on Neurophysiology and Psychophysics* (in the press).
- Colas-Baudelaire, P. (1973) Digital picture processing and psychophysics: a study of brightness perception. Report No. UTEC-CSC-74-025 from the Department of Computer Science, University of Utah.
- Freuder, E. C. (1974) A computer vision system for visual recognition using active knowledge. *M.I.T. A.I. Lab. Technical Report 345*.
- Glass, L. (1969) Moire effect from random dots. *Nature*, 243, 578-580.
- Gilchrist, A. L. (1977) Perceived lightness depends on perceived spatial arrangement. *Science*, 195, 185-187.
- Horn, B. K. P. (1974) Determining lightness from an image. *Computer Graphics and Image Processing*, 3, 277-299.
- Horn, B. K. P. (1977) Image intensity understanding. *Artificial Intelligence*, in the press.
- Hubel, D. H. & Wiesel, T. N. (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol., Lond.* 160 108-154.
- Huffman, D. A. (1971) Impossible objects as nonsense sentences. In *Machine Intelligence 6*, Eds. R. Meltzer & D. Michie, pp 295-323. Edinburgh: The Edinburgh University Press.
- Ittelson, W. H. (1960) *Visual space perception*, pp 145-147. New York: Springer.
- Julesz, B. (1971) *Foundations of Cyclopean perception*. Chicago: The University of Chicago

Press.

Julesz, B. (1975) Experiments in the visual perception of texture. *Scientific American* 232, 34-43, April.

Julesz, B. & Miller, J. E. (1976) Independent spatial-frequency-tuned channels in binocular fusion and rivalry. *Perception* 4, 125-143.

Kolers, P. A. (1972) *Aspects of motion perception*. New York: Pergamon Press.

Land, E. H. & McCann, J. J. (1971) Lightness and retinex theory. *J. opt. Soc. Am.* 61, I-II.

Mackworth, A. K. (1973) Interpreting pictures of polyhedral scenes. *Artificial Intelligence* 4, 121-138.

Marr, D. (1974) A note on the computation of binocular disparity in a symbolic, low-level visual processor. *M.I.T. A.I. Lab. Memo* 327.

Marr, D. (1976) Early processing of visual information. *Phil. Trans. Roy. Soc. B* 275, 483-524.

Marr, D. (1977) Analysis of occluding contour. *Proc. Roy. Soc. B* (in the press).

Marr, D. & Nishihara, H. K. (1977) Representation and recognition of the spatial organization of three-dimensional shapes. (Submitted to the Royal Society for publication).

Marr, D. & Poggio, T. (1976) Cooperative computation of stereo disparity. *Science* 194, 283-287.

Marr, D. & Poggio, T. (1977a) From understanding computation to understanding neural circuitry. *NRP Bulletin on Neurophysiology and Psychophysics*, (in the press).

Marr, D. & Poggio, T. (1977b) A theory of human stereo vision. (In preparation).

Marr, D. , Poggio, T. & Palm, G. (1977c) Analysis of a cooperative stereo algorithm. (In preparation).

- Mayhew, J. E. W. & Frisby, J. P. (1976) Rivalrous texture stereograms. *Nature* 264, 53-56.
- Nevatia, R. (1974) Structured descriptions of complex curved objects for recognition and visual memory. *Stanford A.I. Memo* 250.
- O'Callaghan, J. F. (1974) Computing the perceptual boundaries of dot patterns. *Computer Graphics and Image Processing* 3, 141-162.
- Richards, W. A. (1971) Anomalous stereoscopic depth perception. *J. Opt. Soc Amer.* 61, 410-414.
- Richards, W. A. (1977) Stereopsis with and without monocular cues. *Vision Res.* (in the press).
- Riley, M. (1977) Discriminant functions in early visual processing. (In preparation).
- Rosenfeld, A., Hummel, R. A. & Zucker, S. W. Scene labelling by relaxation operations. *IEEE Transactions on Systems, Man and Cybernetics*, SMC-6, 420-433.
- Schatz, B. R. (1977) On the computation of texture discrimination. *To be given at the Fifth International Joint Conference on Artificial Intelligence, August 1977.*
- Stevens, K. A. (1977) Computation of locally parallel structure. *M.I.T. A.I. Lab Memo* 392.
- Tenenbaum, J. M. & Barrow, H. G. (1976) Experiments in interpretation-guided segmentation. *Stanford Research Institute Technical Note* 123.
- Ullman, S. (1976a) On visual detection of light sources. *Biol. Cybernetics* 21, 205-212.
- Ullman, S. (1976b) Filling-in the gaps: The shape of subjective contours and a model for their generation. *Biol. Cybernetics* 25, 1-8.
- Ullman, S. (1977) The interpretation of visual motion. *M.I.T. Ph. D. Thesis*, June.
- Wallach, H. & O'Connell, D. N. (1959) The kinetic depth effect. *J. exp. Psychol.* 49, 205-217.

Waltz, D. (1975) Understanding line drawings of scenes with shadows. In: *The psychology of computer vision*, Ed. P. H. Winston, pp19-91. New York: McGraw-Hill.